
Local Nonstationarity for Efficient Bayesian Optimization

Ruben Martinez-Cantin
 Centro Universitario de la Defensa
 Zaragoza, 50090, Spain
 rmcantin@unizar.es

Abstract

Bayesian optimization has shown to be a fundamental global optimization algorithm in many applications: ranging from automatic machine learning, robotics, reinforcement learning, experimental design, simulations, etc. The most popular and effective Bayesian optimization relies on a surrogate model in the form of a Gaussian process due to its flexibility to represent a prior over function. However, many algorithms and setups relies on the stationarity assumption of the Gaussian process. In this paper, we present a novel nonstationary strategy for Bayesian optimization that is able to outperform the state of the art in Bayesian optimization both in stationary and nonstationary problems.

1 Introduction

Many problems in engineering, computer science, economics, etc., require to find the extremum of a real valued function. These functions have typically nice properties from a numerical point of view, like being continuous and sometimes smooth (e.g.: Lipschitz continuous). However, many of those functions represent costly processes, expensive trials or several time consuming computations. Furthermore, those functions might not have a closed-form expression or might be highly multimodal.

Bayesian optimization, although being a classic method [13], has become quite popular recently for being a sample efficient method of global optimization [9]. Recent works have found connections with Bayesian optimization and the way biological systems adapt and search, such as human active search [1] or animal adaptation to injuries [3]. In machine learning, it has been applied for automatic algorithm tuning [18] and reinforcement learning [12]. It is specially suitable for these kind of expensive black-box functions and trial-and-error methodologies for using a Bayesian surrogate model, that is, a distribution over target functions $P(f)$. This surrogate model has a twofold purpose:

- It can be updated recursively as outcomes are available from the evaluated trials $\mathbf{y}_i = f(\mathbf{x}_i)$

$$P(f|\mathbf{x}_{1:i}, \mathbf{y}_{1:i}) = \frac{P(\mathbf{x}_i, \mathbf{y}_i|f)P(f|\mathbf{x}_{1:i-1}, \mathbf{y}_{1:i-1})}{P(\mathbf{x}_i, \mathbf{y}_i)}, \quad \forall i = 2 \dots n \quad (1)$$

- It allows a best response analysis for the decision/action \mathbf{a} of selecting the next trial \mathbf{x}_{n+1} :

$$\mathbf{a}^{BO} = \arg \min_{\mathbf{a}} \int_F \delta_n(f, \mathbf{a}) dP(f|\mathbf{x}_{1:n}, \mathbf{y}_{1:n}) \quad (2)$$

where $\delta_n(f, \mathbf{a})$ is the optimality criteria of regret function that drives the optimization towards the optimum \mathbf{x}^* , such as the *optimality gap* $\delta_n(f, \mathbf{a}) = f(\mathbf{x}_n) - f(\mathbf{x}^*)$, the *Euclidean distance error* $\delta_n(f, \mathbf{a}) = \|\mathbf{x}_n - \mathbf{x}^*\|_2$ or the *relative entropy* $\delta_n(f, \mathbf{a}) = H(\mathbf{x}^*|\mathbf{x}_{1:n-1}) - H(\mathbf{x}^*|\mathbf{x}_{1:n})$. The first

condition allows to make the decisions of the second condition with all the information available at that time, thus improving the search of the optimum.

Without loss of generality, consider the problem of finding the minimum of an unknown real valued function $f : \mathbb{X} \rightarrow \mathbb{R}$, where \mathbb{X} is a compact space, $\mathbb{X} \subset \mathbb{R}^d$, $d \geq 1$. For the remainder of the paper, we are going to assume that the surrogate model $P(f)$ is a Gaussian process $\xi(\mathbf{x})$ with inputs $\mathbf{x} \in \mathbb{X}$ and an associate kernel or covariance function $k(\cdot, \cdot)$. One advantage of using Gaussian processes (GPs) as a prior distributions over functions is that new observations of the target function (x_i, y_i) can be easily used to update the distribution over functions. Furthermore, the posterior distribution is also a GP: $\xi_i = [\xi(\cdot)|x_{1:i}, y_{1:i}]$. Therefore, the posterior can be used as an informative prior for the next iteration in a recursive algorithm.

The following algorithm summarizes the steps in Bayesian optimization. Note that, without loss of generality, for the remainder of the paper we are going to assume a Gaussian process with MCMC on the hyperparameters as surrogate model and the expected improvement as the acquisition function. However, the proposed algorithms also work with other popular models such as Student-t processes [14, 17], a discrete representation of the hyperparameters [10] or different acquisition functions such as upper confidence bound [21] or relative entropy [7, 6], among others.

Algorithm 1 Bayesian optimization (BO) with MCMC

```

1: for  $n = 1 \dots N$  do
2:    $\mathbf{X} \leftarrow \mathbf{x}_{1:n-1}$   $\mathbf{y} \leftarrow y_{1:n-1}$ 
3:   for  $i = 1 \dots m$  do ▷ We have  $m$  MCMC samples
4:      $\mu_i \leftarrow \mathbf{k}(\mathbf{x}, \mathbf{X}|\theta_i) \mathbf{K}(\mathbf{X}, \mathbf{X}|\theta_i) \mathbf{y}$  ▷ Predicted mean
5:      $\sigma_i \leftarrow k(\mathbf{x}, \mathbf{x}|\theta_i) - \mathbf{k}(\mathbf{x}, \mathbf{X}|\theta_i) \mathbf{K}(\mathbf{X}, \mathbf{X}|\theta_i)^{-1} \mathbf{k}(\mathbf{X}, \mathbf{x}|\theta_i)$  ▷ Predicted variance
6:   end for
7:    $\Theta \leftarrow \{\theta_i\}_{i=1}^m$  ▷ Update the hyperparameters using slice sampling
8:    $\mathbf{x}_n = \arg \max_{\mathbf{x}_n} \sum_{i=1}^m [(\mathbf{y}_{best} - \mu_i) \Phi(z_i) + \sigma_i \phi(z_i)]$  ▷ Expected improvement
9:    $\mathbf{y}_n \leftarrow f(\mathbf{x}_n)$ 
10: end for

```

The main contribution of the paper is an algorithm for improved Bayesian optimization using a combination of local and global kernels to achieve a nonstationary behavior, called *Spartan Bayesian Optimization*. Although it reaches its best performance in problems that are intrinsically nonstationary, our evaluation shows that it can improve the results of Bayesian optimization in many setups. Additionally, we also present a method to actively learn the Gaussian process hyperparameters while conducting Bayesian optimization with two alternatives: using explicit exploration driven by information gain and using simultaneous perturbations to numerically estimate the variability of the model. Finally, we add some comments about using hierarchical Bayesian optimization to deal with complex input spaces.

2 Nonstationarity in Gaussian processes

Many applications of Gaussian process regression, including Bayesian optimization, are based on the assumption that the process is stationary and, in many times, isotropic. For example, the use of the isotropic squared exponential kernel in GPs is quite frequent: $k_{SE}(\mathbf{x}, \mathbf{x}') = \exp((\mathbf{x} - \mathbf{x}')^T \Lambda (\mathbf{x} - \mathbf{x}'))$, where $\Lambda = \theta_l^{-1} \mathbf{I}$ and θ_l represents the length-scale hyperparameter that captures the smoothness or variability of the function. That is, small values of θ_l will be more suitable to capture signals with high frequency components; while large values of θ_l result in model for low frequency signals or flat functions.

This property results in an interesting behavior in Bayesian optimization. For the same distance between points, a kernel with smaller length-scale will result in higher variance, therefore the exploration will be more aggressive. This idea has been explored previously in [22] by forcing smaller scale parameters to improve the exploration. More formally, in order to achieve no-regret convergence to the minimum, the target function must be an element of the reproducing kernel Hilbert space (RKHS) characterized by the kernel $k(\cdot, \cdot)$ [2, 21]. For a set of kernels like the SE or Matérn, it can be shown that, given two kernels k_l and k_s with large and small length scale hyperparameters respectively, *any function f in the RKHS characterized by a kernel k_l is also an element of the RKHS*

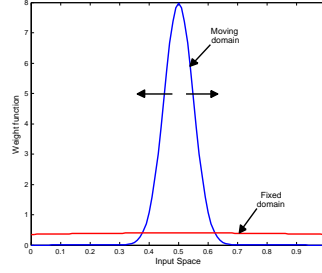


Figure 1: Weighting function of the local and global kernels for nonstationary Bayesian optimization. The global kernel (in red) guarantees full support over the whole input space. The weight of local kernel (in blue) is based on a narrow Gaussian which can be moved through the space, therefore focusing on the region near the minimum value.

characterized by k_s [22]. Thus, using k_s instead of k_l is safer in terms of guaranteeing convergence. However, if the small kernel is used everywhere, it might result in unnecessary sampling of smooth areas.

Nevertheless, most of the applications where Bayesian optimization is used are heterotropic and nonstationary. For example, we may have an heteroscedastic function with different variability, frequency or smoothness in different regions. Thus, these functions require kernels with different length scales for those regions. Take for example a reinforcement learning problem where many policies might result in a failure condition, returning a similar null reward. Therefore, the reward function is almost flat except for a set of parameters where it actually varies.

There has been several attempts to model nonstationary functions with Gaussian processes. The most popular is the use of specific nonstationary kernels [15], Bayesian treed GP models [5] or projecting the input space to a stationary latent space [16]. Recently, a version of the later idea has been applied to Bayesian optimization [19].

Our approach to nonstationarity, the *Spartan Bayesian Optimization* algorithm, is based on the model presented in [10] where the input space is partitioned in different regions such as the resulting GP is the linear combination of local GPs: $\xi(\mathbf{x}) = \sum_i \lambda_i(\mathbf{x}) \xi_i(\mathbf{x})$. Each local GP has its own specific hyperparameters, making the final GP nonstationary even when the local GPs are stationary. In order to achieve smooth interpolation between regions, Krause and Guestrin [10] suggest the use of a weighting function $\nu_i(\mathbf{x})$ for each region, having the maximum in region i and decreasing its value with distance to region i . Then, we can set $\lambda_i(\mathbf{x}) = \sqrt{\frac{\nu_i(\mathbf{x})}{\sum_j \nu_j(\mathbf{x})}}$.

For Bayesian optimization, we suggest the combination of a local and global kernel with multivariate normal distributions as weighting functions as presented in Figure 1. Assuming that the input space is bounded to the $[0, 1]^d$ hypercube, which can be achieved by rescaling the original problem, we consider that for each dimension:

$$\nu_{global}^{(k)} = \mathcal{N}(0.5, 10); \quad \nu_{local}^{(k)} = \mathcal{N}(\theta_{pos}^{(k)}, 0.05) \quad \forall k = 1 \dots d \quad (3)$$

where $\{\theta_{pos}^{(k)}\}_1^d$ is considered to part of the set of hyperparameters of the surrogate model that are learned accordingly when new data is available $\Theta = \{\theta_{pos}, \theta_{local}, \theta_{global}\}$. In that way, the position of the local kernel is adapting towards to the area near the minimum or other important area.

The intuition behind this setup is the same of many adquisition functions in Bayesian optimization: the aim of the surrogate model is not to approximate the target function precisely in every point, but to provide information about the location of the minimum. For example, the resulting model could “flat out” most of the search space, as soon as the region near the minimum have the correct variability. Many optimization problems are difficult due to the fact that the region near the minimum has higher variability that the rest of the space, like the function in Figure 2. However, it is important to note that the kernel hyperparameters are initialized with the same prior for the local and global kernel. Thus, there is guarantee that the local kernel became the kernel with smaller length-scale. Depending on the data captured, it could learn a model where the local kernel has larger length-scale (i.e.: smoother) than the global kernel.

Algorithm 2 Spartan Bayesian Optimization (SBO) with MCMC

```

1: for  $n = 1 \dots N$  do
2:    $\mathbf{X} \leftarrow \mathbf{x}_{1:n-1}; \mathbf{y} \leftarrow y_{1:n-1};$ 
3:   for  $i = 1 \dots m$  do ▷ We have  $m$  MCMC samples
4:      $k(\mathbf{x}, \mathbf{x}' | \theta_i) \leftarrow \lambda_l(\mathbf{x} | \theta_i^{pos}) \lambda_l(\mathbf{x}' | \theta_i^{pos}) k_l(\mathbf{x}, \mathbf{x}' | \theta_i^l) + \lambda_g(\mathbf{x}) \lambda_g(\mathbf{x}') k_g(\mathbf{x}, \mathbf{x}' | \theta_i^g)$ 
5:      $\mu_i \leftarrow \mathbf{k}(\mathbf{x}, \mathbf{X} | \theta_i) \mathbf{K}(\mathbf{X}, \mathbf{X} | \theta_i) \mathbf{y}$  ▷ Predicted mean
6:      $\sigma_i \leftarrow k(\mathbf{x}, \mathbf{x} | \theta_i) - \mathbf{k}(\mathbf{x}, \mathbf{X} | \theta_i) \mathbf{K}(\mathbf{X}, \mathbf{X} | \theta_i) \mathbf{k}(\mathbf{X}, \mathbf{x} | \theta_i)$  ▷ Predicted variance
7:   end for
8:    $\Theta \leftarrow \{\theta_{global}^{(i)}, \theta_{local}^{(i)}, \theta_{pos}^{(i)}\}_{i=1}^m$  ▷ Update the hyperparameters using slice sampling
9:    $\mathbf{x}_n = \arg \max_{\mathbf{x}_n} \sum_{i=1}^m [(\mathbf{y}_{best} - \mu_i) \Phi(z_i) + \sigma_i \phi(z_i)]$  ▷ Expected improvement
10:   $\mathbf{y}_n \leftarrow f(\mathbf{x}_n)$ 
11: end for

```

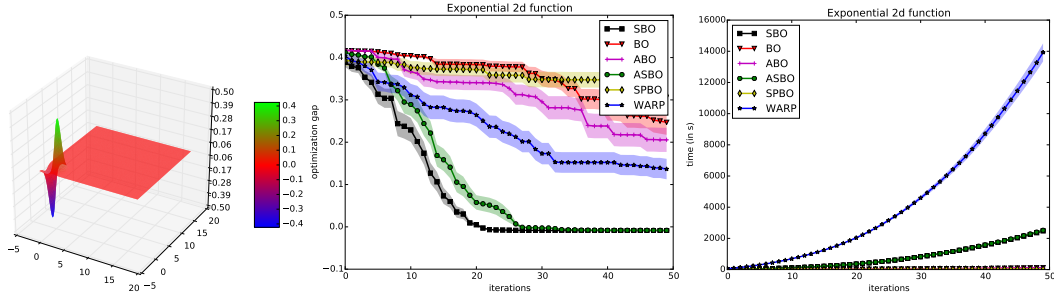


Figure 2: Results for the exponential 2D function. Center: we can see the optimization gap. Nonstationary methods, such as *WARP* and the proposed method *SBO* result in an improved convergence. The suggested method was able to find the minimum always in less than 20 iterations. Right: Total optimization time in seconds.

3 Active hyperparameter learning during Bayesian optimization

As commented previously, learning accurate kernel hyperparameters is of paramount importance in Bayesian optimization, specially in the case of nonstationary models, where the number of hyperparameters increases considerably. On the other hand, Bayesian optimization algorithms are rooted on the idea of having few samples. Therefore, learning those hyperparameters becomes a challenging problem.

In this section, we present a method to actively learn the kernel hyperparameters while doing optimization, by explicitly exploring areas with high information gain over the hyperparameters.

3.1 Active hyperparameter learning through Information Gain

In order to improve the quality of the surrogate model, we will try to reduce the entropy of the hyperparameters $H(\Theta)$. Implicitly, this is done already in Bayesian optimization. Many popular criteria such as the expected improvement, the upper confidence bound, etc., include an exploration term which tries to minimize the predicting variance of the surrogate model. The *information never hurts* principle shows that any exploration strategy will, in expectation, decrease $H(\Theta)$. See Proposition 8 of [10] for details.

Krause and Guestrin [10] also suggest a criterion for explicit exploration based on the information gain (IG) of the GP hyperparameters. Following that criterion, the decision for the next point to evaluate will be:

$$\begin{aligned}
 c_{IG} &= H(\Theta | \mathbf{y}_{1:n}) - H(\Theta | \mathbf{y}_{1:n}, \mathbf{y}_{n+1}) = \arg \max_{\mathbf{x}_{n+1}} H(\mathbf{y}_{n+1} | \mathbf{y}_{1:n}) - H(\mathbf{y}_{n+1} | \mathbf{y}_{1:n}, \Theta) \\
 &\approx -\log \left\{ \sum_{i=1}^m \left[(\mu_i - \hat{\mu})^2 + \sigma_i^2 \right] \right\} + \sum_{i=1}^m \log \sigma_i^2
 \end{aligned} \tag{4}$$

where $\hat{\mu} = \sum_i \mu_i / m$. Any other sampling distribution could be used for Θ adding the corresponding weight to each sample. For example, in the original work [10], a discrete distribution is used.

The problem with \mathbf{a}_{IG} is that it is purely exploratory criterion, which is not intended for optimization. Thus, it needs to be combined with other optimization criterion. We suggest a linear combination with an annealing coefficient to gather information about the hyperparameters Θ at the beginning, and focusing on finding the minimum \mathbf{x}^* after several iterations. For example, the expected improvement with explicit information gain criterion is defined as:

$$c_{EIIIG} = \sum_{i=1}^m [(y_{best} - \mu_i) \Phi(z_i) + \sigma_i \phi(z_i)] + \frac{\alpha}{n^2} \left(\log \left\{ \sum_{i=1}^m [(\mu_i - \hat{\mu})^2 + \sigma_i^2] \right\} + \sum_{i=1}^m \log \sigma_i^2 \right) \quad (5)$$

where $z_i = (y_{best} - \mu_i) / \sigma_i$. The functions $\Phi(\cdot)$ and $\phi(\cdot)$ are the normal cdf and pdf respectively. The coefficient α/n^2 represents the importance information gain component with respect to the expected improvement. Also, the coefficient is annealed to increase the effect of the IG early on to help learning the kernel hyperparameters. Eventually $\alpha/n^2 \rightarrow 0$, resulting in the classical expected improvement once the hyperparameters are good enough.

3.2 Simultaneous perturbation Bayesian optimization

Analyzing the shape of the information gain criterion we found that, in many cases, the highest value was near a previous sample. We concluded that the leght-scale hyperparameter is related to the variability of the function, which, at the same time, it is related to the local gradient. Therefore, in a certain way, the information gain criterion was estimating the gradient numerically to obtain information of the variability of the function. It is known that adding gradient information can improve considerably the accuracy of Gaussian process regression [15]. However, in many applications, the gradient is not available. Furthermore, using finite difference methods require many samples, which is against the philosophy of Bayesian optimization.

In the field of stochastic optimization, the simultaneous perturbation stochastic approximation algorithm (SPSA) [20] is a variation of the classic finite difference algorithm for high dimensional spaces. In this method, the gradient is approximated by finite differences of a small subset of randomly sampled perturbations. The beauty of the SPSA algorithm is that, although the gradient is not perfectly estimated, the algorithm is guaranteed to converge to the local optimum.

Algorithm 3 Simultaneous Perturbation Bayesian Optimization (SPBO)

```

1: for  $i = 1 \dots N/2$  do ▷ Generates two samples per iteration
2:    $\mathbf{x}_i \leftarrow$  Computed with Algorithm 1.
3:   if  $i < T$  then ▷ We compute perturbations only at the beginning, during exploration.
4:      $\Delta_x$  perturbation vector. Each component sampled from a  $\pm 1$  Bernoulli distribution
5:      $\mathbf{x}_j \leftarrow \mathbf{x}_i + \frac{c}{i^\gamma} \Delta_x$ 
6:   end if
7: end for
```

We present an algorithm called Simultaneous Perturbation Bayesian Optimization, which, for every sample using Bayesian optimization, also computes a perturbed sample. Theoretically, this algorithm reduce the computational burden of Bayesian optimization and information gain inasmuch as it is applied half of the iterations. Compared to Bayesian optimization, the cost of sampling a random perturbation is negligible with respect to the cost of updating the Gaussian process model and computing the maximum information gain. Also, note how the perturbation can be computed only in the first part of the optimization process, similar to the annealing process in the previous section.

4 Hierarchical Bayesian optimization

In many Bayesian optimization applications, it is becoming a common trend to simultaneously optimize different kinds of variables, for example: continuous, discrete, categorical, etc. While Gaussian processes are suitable for modeling those spaces, Bayesian optimization can become quite involved in line 8 of Algorithm 1, as the acquisition function must be optimized in the same input space.

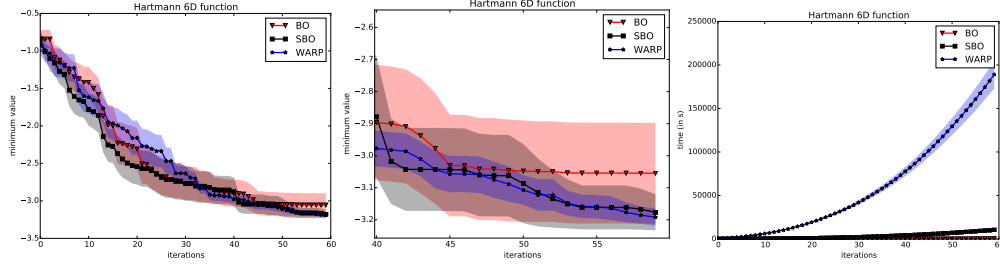


Figure 3: Hartmann 6D function. Left: Minimum value obtained. Center: zoom of the minimum value in the last iterations. Right: time in seconds.

Available implementations of Bayesian optimization like Spearmint [18] use grid sampling and rounding tricks to combine different input spaces. However, this might reduce the quality of the final result [11] compared to proper optimization methods.

In this section, we suggest to use a hierarchical Bayesian optimization model, where the input space is partitioned between homogeneous variables, for example: continuous variables $\mathbf{x}^{(c)}$ and discrete variables $\mathbf{x}^{(d)}$. Therefore, the evaluation of an element higher in the hierarchy implies the full optimization of the elements lower in the hierarchy. In principle, that would require much more function evaluations but, as the input space has been partitioned, the dimensionality of each separate problem is much lower. In practice, for the same number of function evaluations, the computational cost is considerably reduced.

Algorithm 4 Hierarchical Bayesian Optimization (HBO) — Total budget $N = N_1 \cdot N_2$

```

1:  $\mathbf{x} = [\mathbf{x}^{(c)}, \mathbf{x}^{(d)}]$  ▷ We separate the discrete and continuous components
2: for  $n = 1 \dots N_c$  do ▷ Outer loop - Continuous optimization
3:   Update model and kernels hyperparameters for  $\mathbf{x}^{(c)}$ 
4:   Find continuous component of next point  $\mathbf{x}_n^{(c)}$ 
5:   for  $k = 1 \dots N_d$  do ▷ Inner loop - Discrete optimization
6:     Update model and kernels hyperparameters for  $\mathbf{x}^{(d)}$ 
7:     Find discrete component of next point  $\mathbf{x}_k^{(d)}$ 
8:      $\mathbf{y}_k \leftarrow f([\mathbf{x}_n^{(c)}, \mathbf{x}_k^{(d)}])$  ▷ Update  $\mathbf{x}_n^{(d)}$ 
9:   end for
10:   $\mathbf{y}_n \leftarrow f([\mathbf{x}_n^{(c)}, \mathbf{x}_n^{(d)}])$  ▷ Update  $\mathbf{x}^{(c)}$  and save corresponding  $\mathbf{x}^{(d)}$ 
11: end for
12:  $\mathbf{x}^* = [\mathbf{x}^{*(c)}, \mathbf{x}^{*(d)}]$ 

```

An advantage of this approach is that we can combine different algorithms for different levels of the hierarchy. For example, using Random Forests [8] could be more suitable as a surrogate model for certain discrete/categorical variables than Gaussian processes. In contrast, we lose the correlation among variables in the inner loop, which might be counterproductive in certain situations.

5 Evaluation and results

For the evaluation of the suggested algorithms we have implemented them using the popular BayesOpt¹ library [11]. This allowed us to compare our proposal with a large variety of surrogate models, kernels, etc. For example, the results presented in this sections are based on the standard convention in Bayesian optimization literature, that is, a simple zero-mean Gaussian process, a Matérn kernel 5/2 with automatic relevant determination for continuous variables, a Hamming kernel as presented in [22] for categorical variables and slice sampling for learning the model hyperparameters (kernel, warping, etc.). However, the suggested method has also been tested with other models such as Student-t processes, other kernels, etc. Due to the computational burden of MCMC

¹The code will be available as GPL once the paper gets published.

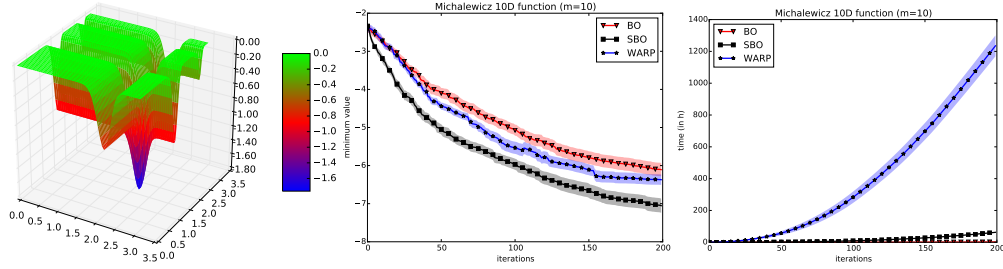


Figure 4: Michalewicz 10D function with $m=10$. On the left, there is a visualization of the first and second dimensions to appreciate the difficulty of this function. The number of local minima increases exponentially with the number of dimensions. Note that the time is in hours.

for the hyperparameters, we have used a small number of samples (10), while trying to decorrelate every resample by large burn-in periods (100 samples) following the convention in [18].

We compare standard Bayesian Optimization (BO), Bayesian Optimization with Active hyperparameter learning (ABO), Spartan Bayesian Optimization (SBO), Spartan Bayesian Optimization with Active hyperparameter learning (ASBO), Simultaneous Perturbation Bayesian Optimization (SPBO). We also implemented the input warping (WARP) of Snoek et al. [19] which, to the authors knowledge, is currently the only Bayesian optimization algorithm specific for nonstationary functions. All experiments were repeated between 10 and 25 times, depending on the problem, using common random number to reduce the sampling error between algorithms. All the experiments include 5-10 initial samples generated through latin hypercube sampling, not included in the plot.

5.1 Optimization benchmarks

We have evaluated the algorithms on a set of standard functions for global optimization. Although popular in the literature, we have avoided the use of the Branin-Hoo and Camelback functions as there is barely room from improvement using “vanilla” BayesOpt [11].

Figure 2 shows the results of optimizing the exponential 2D function $f(\mathbf{x}) = x_1 \exp(-x_1^2 - x_2)$ for $x_1, x_2 \in [-2, 18]^2$ from [5]. The use of classical stationary models (BO) results in a poor convergence because of the high nonstationarity of the function. Using an active strategy speed up the results, requiring slightly fewer iterations. Clearly, nonstationary methods, such as Snoek et al [19] (WARP) and the proposed method (SBO) result in an improved convergence. In fact, the convergence of the proposed method requires a number of iterations so small that actively learning the kernel hyperparameters results in a waste of samples.

The computation times of the different algorithms are clearly driven by the dimensionality and shape of the distribution of the MCMC. The cost of the active component is negligible when compared to their passive counterpart. We found that learning the Beta parameters for the warping function was extremely slow as many samples were rejected. This could be alleviated with a different MCMC algorithm such as hybrid Monte Carlo or sequential Monte Carlo samplers. However, note that we use slice sampling as recommended by the authors [19].

For the Hardmann 6D function (shown in Figure 3), the differences are not statistically significant, which shows that when the function is stationary, nonstationary methods are no worse than standard Bayesian optimization. Furthermore, we can see that at the end they are more robust.

The Michalewicz function is known to be one of the hardest benchmarks in surrogate based global optimization. Figure 4 shows the results. The active algorithms as well as the SPBO have been removed as they were statistically indifferent from their passive counterpart.

5.2 Surrogate benchmarks for machine learning problems

Our next set of experiments is based on well known benchmarks for automatic tuning of machine learning problems. However, in order to simplify the evaluation, we have used the surrogate benchmarks presented in [4]. Among all the available benchmarks we have selected the Gradient Boosting

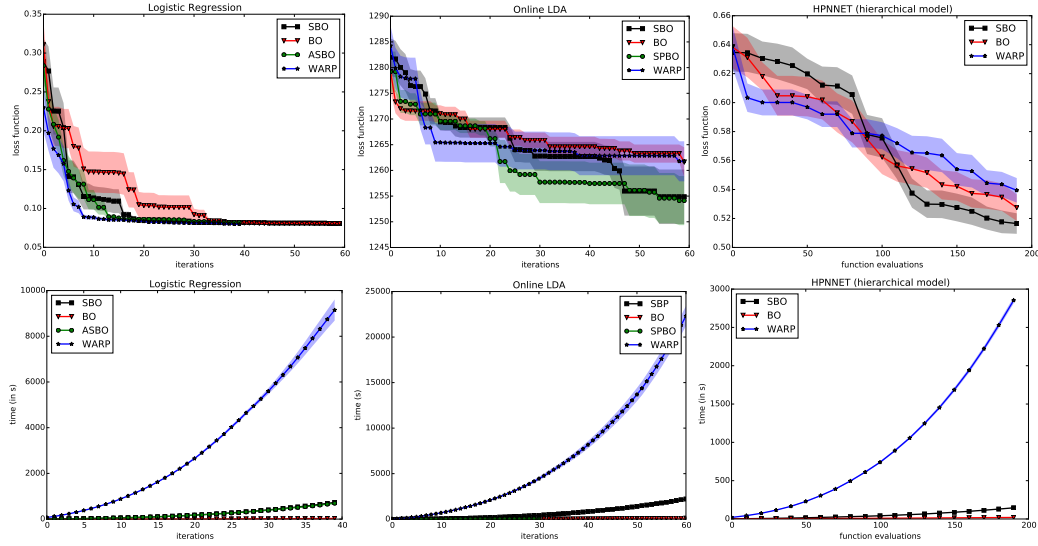


Figure 5: Surrogate benchmarks [4] based on real hyperparameter optimization of machine learning algorithms. From left to right: logistic regression (4D continuous), online LDA (3D continuous) and deep neural network (HP-NNET with the mrbi dataset, 7D continuous, 7D categorical). Top row: loss functions. Bottom row: computational time, in seconds.

as it provides the lowest RMSE with respect to the actual problem. We explicitly avoid the Gaussian process surrogate to avoid the advantage of perfectly modeling. The results are shown in Figure 5. For clarity, we show only the statistically significant results apart from BO, local and warp.

First, the logistic regression is an easy problem for Bayesian optimization. Even the simple Bayesian optimization can reach the minimum in less than 40 iterations. In this case, the warped method is the fastest one, with less than 10 iterations. However, the proposed method is comparable specially when using active learning, by a fraction of the total cost. It is important to note that, although Bayesian optimization is intended for expensive functions and the cost per iteration is negligible, we are talking of minutes to hours for single iteration of the warped method in a standard laptop. For the onlineLDA problem, the warped method gets stuck like the standard Bayesian optimization, while our method is able to escape from the local minima. Surprisingly, the SPBO method is the best in this example, even though the computational cost is the lowest by a large margin. This shows that for certain applications and structured problems, gradient information is fundamental.

Finally, we evaluate the HP-NNET problem. In this case, due to the high dimensionality and heterogeneity of the input space (7 continuous + 7 categorical parameters) we have applied the hierarchical model presented in Section 4. Also, in order to reduce the computational cost, the nonstationary (warping or local kernel) is applied only on the continuous variables (outer loop). Note that, in this case, the plots are with respect to target function evaluations. In this case, due to the complexity of the problem, the local kernel fails to converge at an early stage. However, with more data available, the local kernel jumps to a good spot and the convergence is faster.

6 Conclusions

We have presented a new algorithm called Spartan Bayesian Optimization which combines a local and a global kernel to deal with nonstationarity in Bayesian optimization. Besides, we have shown that the model can improve convergence speed even in stationary problems. The suggested algorithm achieves similar or better results than the state of the art by a small fraction of the computational cost. In addition to that, we have presented some ideas to improve the results of any Bayesian optimization algorithm: by actively learning the surrogate hyperparameters, by efficiently estimating the gradient of the target function and by building a hierarchical model to reduce the heterogeneity in the input space.

References

- [1] Ali Borji and Laurent Itti. Bayesian optimization explains human active search. In C.J.C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 55–63. Curran Associates, Inc., 2013.
- [2] Adam D. Bull. Convergence rates of efficient global optimization algorithms. *Journal of Machine Learning Research*, 12:2879–2904, 2011.
- [3] Antoine Cully, Jeff Clune, Danesh Tarapore, and Jean-Baptiste Mouret. Robots that can adapt like animals. *Nature*, 521:503507, 2015.
- [4] K. Eggenberger, F. Hutter, H.H. Hoos, and K. Leyton-Brown. Efficient benchmarking of hyperparameter optimizers via surrogates. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, January 2015.
- [5] Robert B Gramacy. *Bayesian treed Gaussian process models*. PhD thesis, University of California, Santa Clara, 2005.
- [6] Philipp Hennig and Christian J. Schuler. Entropy search for information-efficient global optimization. *Journal of Machine Learning Research*, 13:1809–1837, 2012.
- [7] José Miguel Hernández-Lobato, Matthew W Hoffman, and Zoubin Ghahramani. Predictive entropy search for efficient global optimization of black-box functions. In Z. Ghahramani, M. Welling, C. Cortes, N.D. Lawrence, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 918–926. Curran Associates, Inc., 2014.
- [8] Frank Hutter, Holger H. Hoos, and Kevin Leyton-Brown. Sequential model-based optimization for general algorithm configuration. In *LION-5*, page 507523, 2011.
- [9] Donald R. Jones, Matthias Schonlau, and William J. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13(4):455–492, 1998.
- [10] Andreas Krause and Carlos Guestrin. Nonmyopic active learning of gaussian processes: an exploration-exploitation approach. In *International Conference on Machine Learning (ICML)*, Corvallis, Oregon, June 2007.
- [11] Ruben Martinez-Cantin. BayesOpt: A Bayesian optimization library for nonlinear optimization, experimental design and bandits. *Journal of Machine Learning Research*, 15:3735–3739, 2014.
- [12] Ruben Martinez-Cantin, Nando de Freitas, Eric Brochu, Jose Castellanos, and Arnaud Doucet. A Bayesian exploration-exploitation approach for optimal online sensing and planning with a visually guided mobile robot. *Autonomous Robots*, 27(3):93–103, 2009.
- [13] Jonas Mockus. *Bayesian Approach to Global Optimization*, volume 37 of *Mathematics and Its Applications*. Kluwer Academic Publishers, 1989.
- [14] Anthony O’Hagan. Some Bayesian numerical analysis. *Bayesian Statistics*, 4:345–363, 1992.
- [15] Carl E. Rasmussen and Christopher K.I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, Cambridge, Massachusetts, 2006.
- [16] Paul D Sampson and Peter Guttorp. Nonparametric estimation of nonstationary spatial covariance structure. *Journal of the American Statistical Association*, 87(417):108–119, 1992.
- [17] Amar Shah, Andrew Gordon Wilson, and Zoubin Ghahramani. Student-t processes as alternatives to Gaussian processes. In *AISTATS, JMLR Proceedings. JMLR.org*, 2014.
- [18] Jasper Snoek, Hugo Larochelle, and Ryan Adams. Practical Bayesian optimization of machine learning algorithms. In *NIPS*, pages 2960–2968, 2012.
- [19] Jasper Snoek, Kevin Swersky, Richard S. Zemel, and Ryan Prescott Adams. Input warping for bayesian optimization of non-stationary functions. In *International Conference on Machine Learning*, 2014.
- [20] J.C. Spall. Implementation of the simultaneous perturbation algorithm for stochastic optimization. *Aerospace and Electronic Systems, IEEE Transactions on*, 34(3):817–823, Jul 1998.
- [21] Niranjana Srinivas, Andreas Krause, Sham Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In *Proc. International Conference on Machine Learning (ICML)*, 2010.
- [22] Ziyu Wang, Masrour Zoghi, Frank Hutter, David Matheson, and Nando de Freitas. Bayesian optimization in high dimensions via random embeddings. In *International Joint Conferences on Artificial Intelligence (IJCAI) - Distinguished Paper Award - Extended version: <http://arxiv.org/abs/1301.1942>*, 2013.